Part 1

Overview

- **Chapter 1: Proposing Virtualization**
- **Chapter 2: The Architecture of Hyper-V**

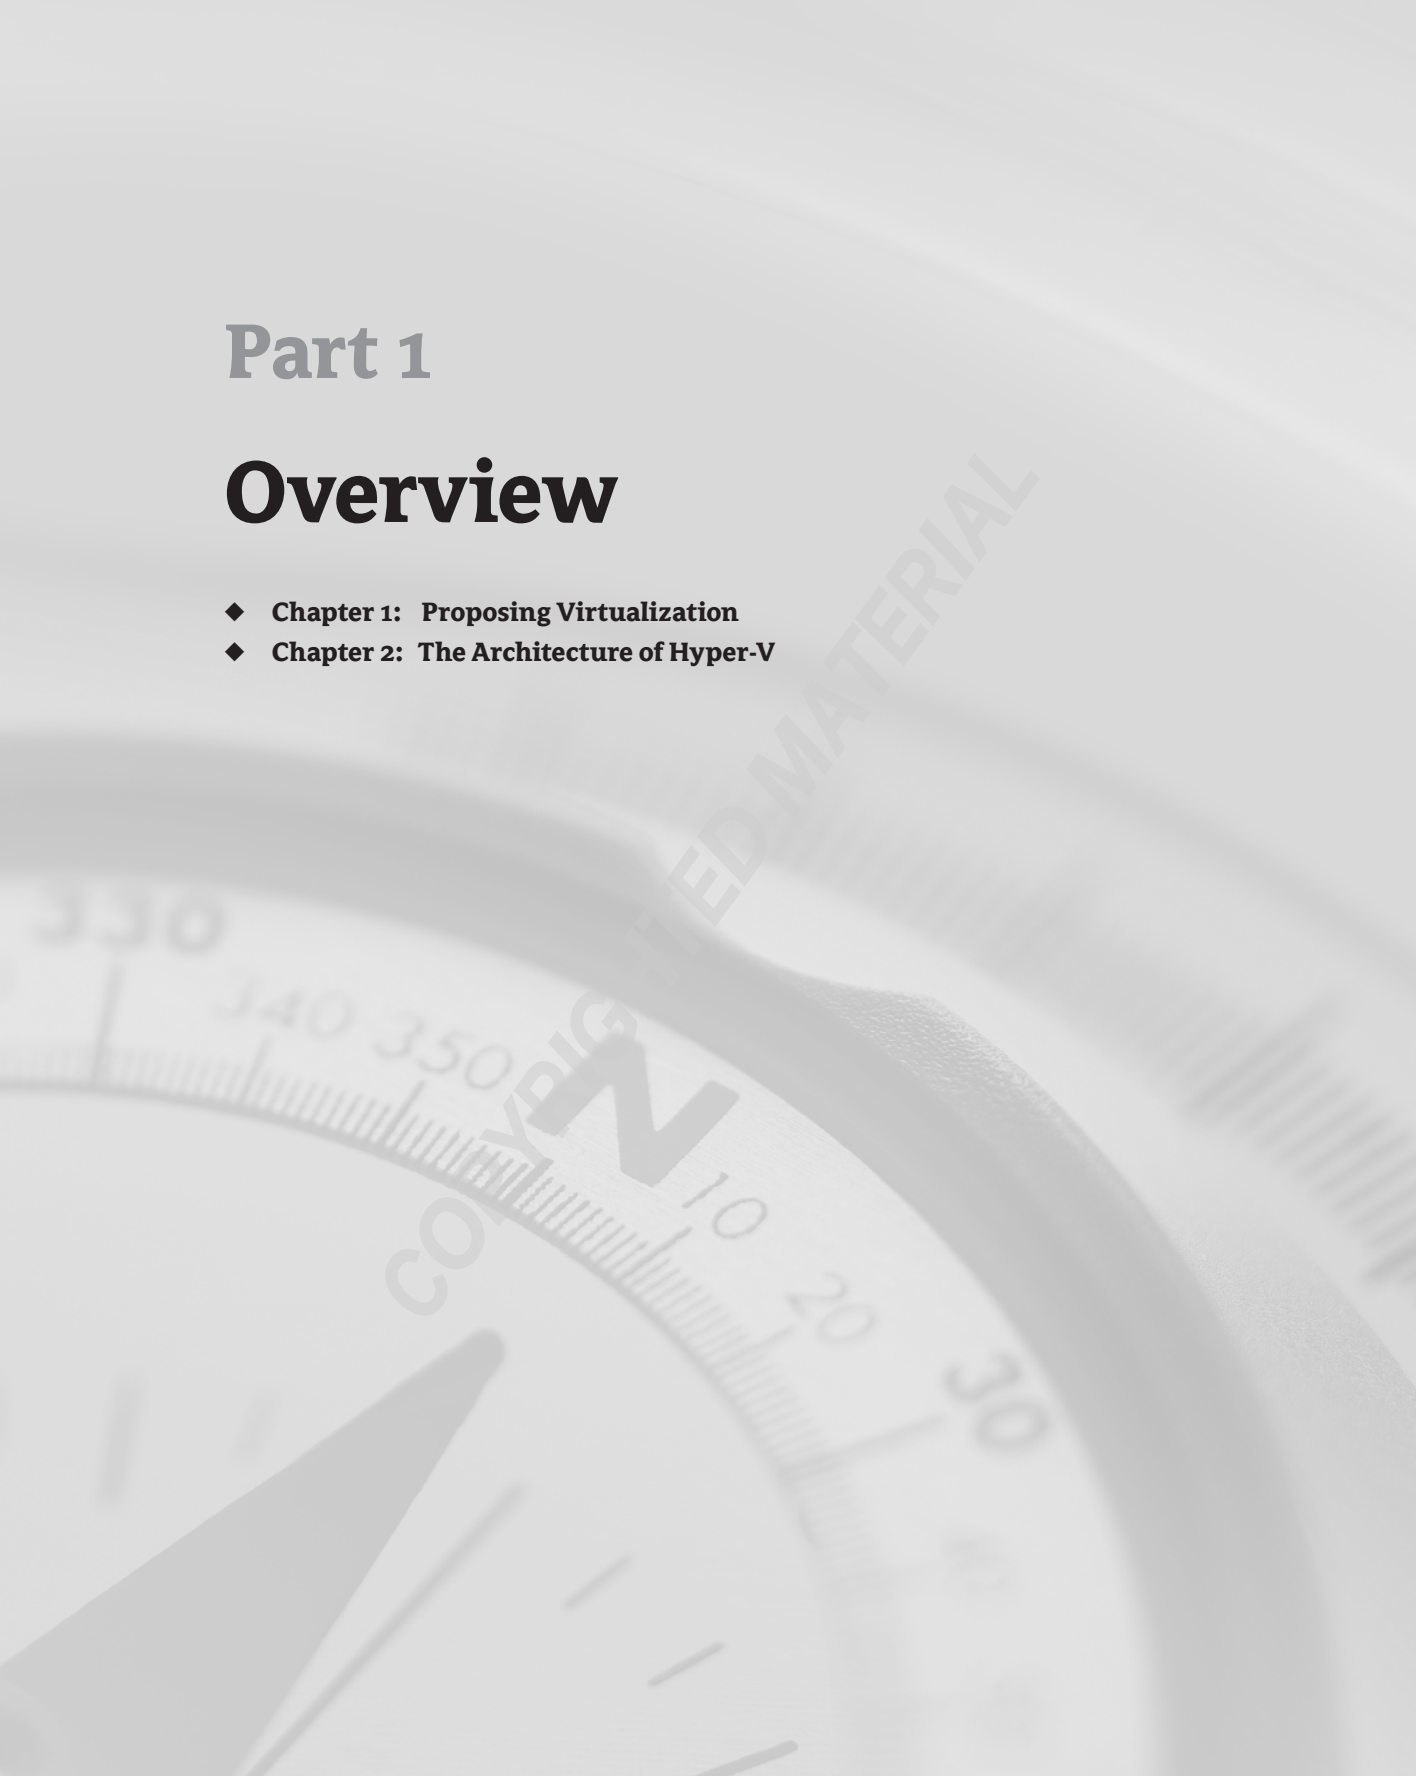

Chapter 1

Proposing Virtualization

The first phase of any project, whether you're an internal consultant or an external consultant, is to develop a proposal for that project so that you can gain executive commitment. The best way to secure acceptance is to cater to both the technical requirements and the often overlooked business case requirements.

The technical solution should be a business-focused solution and be understandable by both technical and nontechnical staff. It should answer the typical questions. What does this project mean for my company? What are the implications for my company? How can we successfully move forward based on our current environment?

The results should be a vision for the project, executive consensus, and a defined set of clear next steps. The proposal should model areas of business value and map technology against that business value. Case studies, reference sites, white papers, and ROI analysis have a part to play when developing both your business case and your technical case.

Interestingly, it appears that Hyper-V is being implemented not for the expected "virtualization project" but as a component of some other solution required by the business.

In this chapter, you will learn what the elements are for both the technical case and the business case when proposing virtualization.

The Business Case

Virtualization enables you to pool and share IT resources to better serve your business and create a business-ready dynamic infrastructure. From a business perspective, the pooling and sharing of IT resources allows IT supply to keep pace with the ever-changeable business demand.

From a cost perspective, the pooling and sharing of resources helps you increase the utilization of IT assets and thus reduce your total cost of ownership (TCO), in terms of both capital expense (CAPEX) and operating expense (OPEX) costs. This is achieved in part by enabling greater use of your physical resources.

Line-of-Business Application Continuity

Customers are turning to virtualization to enhance their existing business continuity strategies (BCSs) and to provide a simplified form of business continuity for existing x86 platforms as they adopt virtualization. Hyper-V provides a robust, reliable, and secure platform that isolates applications and operating systems from their underlying hardware, dramatically reducing the complexity of implementing and testing business continuity service.

In its most simplistic form, this involves the implementation of replicated storage to support the constituent parts of the virtual machine. Storage vendors for the most part provide either in-box or add-on replication capabilities, which are easily enabled. Replicating the storage presented to the virtualized infrastructure, even without storage array–based replication techniques leveraging software vendors, provides the basis for a business continuity service.

As long as there is sufficient capacity at the designated business continuity site, the virtual machines being protected — independent of the underlying servers, storage infrastructure, and networking — allow for the quantity of servers at the business continuity site to be different. This is in contrast to a traditional x86 business continuity solution, which typically involved maintaining a direct one-to-one relationship between the production and the business continuity site, in terms of servers, storage infrastructure, management infrastructure, and networking hardware.

Storage replication is a simple, yet powerful, approach. However, there are a number of considerations to be made to implement this type of solution in an effective manner. Building a generic business continuity solution can be extremely complex and most physical and virtual implementations, while often automated, are often heavily customized to meet both business and technical requirements.

Centralized Computing

Virtualization offers new methods of implementing centralized computing. Take the concept of a Virtual Desktop Infrastructure (VDI). VDI introduces a new way of managing end-user computing environments. VDI allows IT administrators to host and administer end-user desktops on a virtualized infrastructure running in the datacenter.

Users access their desktop using a remote desktop protocol and often a thin client device. While sharing similarities with other computing models, VDI offers many new and compelling benefits for increasing the manageability, performance, and security of end-user devices. Although VDI is not the panacea today, it is an architecture that is leveraging a centralized computing model.

Centralized computing, however, is not limited to virtualization. Centralized computing through physical consolidation is where servers, storage, and networking from many locations, typically datacenters, are physically relocated to fewer locations. The IT infrastructure is brought under a common operational framework, which has the following advantages:

- Consistent level of service
- Improved level of security
- Reduced operational costs
- Standardized management approach
- Clearer understanding of maintenance, power, and cooling costs

Conversely, centralized computing through virtualization is where the hardware remains in the same location but the number of underutilized servers is reduced, or consolidated, using a consistent methodology to map physical assets to virtual ones. The advantages of this type of centralized computing model are similar to physical consolidation:

- Consistent level of service
- Reduced operational costs
- Standardized management approach

Lower Costs

Virtualization helps drive down both CAPEX and OPEX costs. CAPEX cost savings can come in the form of reduced expenses for the acquisition of hardware and the potential savings for datacenter real estate. OPEX costs stem from the reduction in power and cooling costs, management costs, and the costs associated with server downtime or outage costs.

To assist you with developing your business case, Microsoft has made available a return on investment (ROI) calculator that was developed independently by leading TCO/ROI experts at Alinean (`www.alinean.com`). The ROI calculator exists as a sales enablement tool and was designed to help quantify the TCO/ROI savings and competitive advantage of Microsoft's integrated virtualization solution. For more details on the Microsoft Integrated Virtualization ROI Calculator, take a look at Chapter 4.

Server Hardware

Moving to a virtual environment can help you cut costs by reducing the number of physical servers necessary to support your infrastructure. By consolidating your server hardware, you will achieve higher utilization levels and thus reduce your overall hardware requirements.

Power and Cooling

Virtualization can help you take control of rapidly rising power and cooling costs. The savings typically stem from reductions in the number of physical servers in your environment. Advances in both modern hardware and operating systems have parts to play. Microsoft introduced the concept of core parking in Windows 2008 R2; core parking is the ability for the operating system to put cores of processors and entire processors into a low power state when not in use.

Hyper-V R2 supports core parking by allowing virtual machine threads to be moved between cores to enable core parking to happen.

Server Provisioning Costs

Virtualization allows you to rapidly provision virtual machines in less time, which in turn leads to reduced infrastructure management costs. Through the use of System Center Virtual Machine Manager (SCVMM), the act of creating a new virtual machine can occur in one of three ways:

- Creating virtual machines from a template
- Using an existing virtual machine (cloning)
- Using a blank virtual hard disk

In addition, SCVMM supports various different options for provisioning virtual machines. By leveraging its built-in PowerShell script library, you can automate the entire provisioning process.

Green Computing

Energy consumption is a critical issue for IT organizations today, whether the goal is to reduce your costs, save the planet, or keep your datacenter running. One of the easiest ways to reduce the energy demands of your datacenter is through server consolidation and dynamic systems management of your server assets.

Virtualization is a fundamental component of a green computing initiative. Consolidating physical servers into one or more virtual servers is a more efficient use of resources, which in turn means less hardware is deployed, and thus less power is consumed. Less hardware and a reduction on overall power consumption means cost savings and a smaller carbon footprint.

However, green computing is not just about consolidating physical servers and what consolidation ratios you achieve. It is as much about the design and manufacture of that hardware, including any peripherals and the entire life cycle of that hardware, from initial purchase to final disposal.

There are a number of industry initiatives. The Green Grid is a global consortium of IT companies seeking to improve energy efficiency within datacenters and business computing environments around the world. Its website (`www.thegreengrid.org`) has a number of useful resources and tools that may assist with the further development of your green computing strategy.

To calculate your potential green savings, Microsoft has developed a sustainability calculator to estimate your company's carbon footprint and see the potential savings that are achievable through consolidation. For a preview of this tool and to see how it could benefit your organization, visit this site:

```
http://www.microsoft.com/environment/greenit/Preview.aspx?type=server
```

Self-Provisioning

SCVMM provides a web-based portal where authorized users can provision new virtual machines without directly involving IT staff. This capability especially targets software test and development teams, which often set up and tear down temporary virtual machines for application development purposes.

SCVMM administrators grant users permissions to create and operate their own virtual machines within a controlled environment and on a limited set of Hyper-V hosts. This limited set of Hyper-V hosts is typically organized within a Host Group, which is a logical container within SCVMM. The SCVMM administrator is required to create self-service user roles, which determine the scope of the user's actions on their designated virtual machines.

To create, operate, and manage their virtual machines, self-service users use the SCVMM Self-Service Portal. This website provides a controlled environment for users in the self-service user role. The administrator determines which host groups the self-service users can create virtual machines on. When a self-service user creates a virtual machine, the virtual machine is automatically placed on the most suitable host in the host group.

Active Directory users or groups can be added to self-service user roles. The permissions granted to the user role apply to all members of the user group whether they are individuals or groups.

Virtual machine owners can be individual users or groups. Under individual ownership, an individual owns, operates, and manages its own virtual machines. Under group ownership, virtual machines are owned, operated, and managed by the group.

You can set a virtual machine quota in a self-service user role to limit the number of virtual machines that a user or group can deploy. Quota points are assigned to the virtual machine template or templates that self-service users use to create their virtual machines.

Quota points apply only to virtual machines on a host. If a self-service user is allowed to store virtual machines, the quota does not apply to virtual machines stored in the SCVMM library. When the self-service user's quota is reached, the user cannot create any new virtual machines until an existing virtual machine is removed or stored.

Business Continuity Planning

Business continuity planning (BCP) is the ability to minimize scheduled and unscheduled downtime, using the host-based failover features of the virtualization platform and guest-based clustering.

Windows 2008 R2 and Hyper-V includes support for host-based clustering of virtual machines. This allows an organization to meet the availability thresholds previously reserved only for cluster-aware applications. Because virtualization clustering allows a guest machine to be transferred across physical nodes with zero downtime, the number of machines that can be targeted for virtualization dramatically increases. This is particularly compelling for applications that grow into mission-critical status but were never designed with high availability in mind.

In addition to the benefits of failover clustering, virtualization of target machines can greatly enhance the business continuity and recovery processes. Because each virtual machine is a collection of files on a physical host, the files (VHDs, AVHDs, and so on) can be moved to a new location, including alternate datacenters, and be brought back online without requiring a complete rebuild. Alternate datacenters can consist of far fewer physical machines, specifically designed to provide an emergency level of service. Windows 2008 R2 and Hyper-V also introduces the feature of Live Migration, where running virtual machines can transfer from one clustered host node to another with zero downtime. There are two options for using Live Migration:

- Planned
- Unplanned

Unplanned migrations occur when the active node (Hyper-V host) running the guests becomes unexpectedly unavailable. In this case, the other nodes in the cluster recognize the failure, cluster resources are moved, and the guest machines are brought back up on available cluster nodes, reducing the overall downtime. Note that in an unplanned migration, the guest machine state is lost and is restarted, just as if the "power cord" were pulled.

In a planned Live Migration, guests are transferred between nodes while they are running. This process is carried out automatically by either the Failover Cluster Management tool, a PowerShell script, or SCVMM. In the BCP scenario, the flexibility of moving machines between

locations supports the business driver for functionality in the event of a disaster. Windows 2008 R2 and Hyper-V provides an organization with the following benefits:

- Ensures continuity or uninterrupted provisioning of operations, servers, and services
- Reduces service interruptions with failover clustering on the host
- Allows almost immediate rebalancing of resources to guest machines to meet growing or changing business requirements
- Improves disaster response and business recovery requirements by minimizing the infrastructure necessary to run all mission-critical applications under a recovery scenario

The Technical Case

A virtualized environment increases IT flexibility because a varied range of resources can be added, changed, or moved as needed to meet the changes in business demand. Resources can be scaled to either increase or decrease based on changing workloads and patterns. As a by-product, virtualization also improves IT's level of resiliency by simplifying backup, failover, disaster recovery, and business continuity solutions.

Test and Development

Virtualization can maximize test hardware, reducing costs, improving life-cycle management, and improving test coverage. Nearly all test and development machines are good candidates for virtualization, unless you are performing specific production workload tests.

Virtualization of the test and development environment allows you to do the following:

- Consolidate hardware resources and better utilize hardware with consolidated workloads.
- Improve and maximize hardware utilization, especially for machines with short lives or involved in destructive life cycles.
- Reduce TCO for hardware, electrical, cooling, power, and rack space footprint.
- Greatly reduce time to provision new servers (on virtual hardware) to developers as fully imaged servers.
- Reduce time to migrate new software from development to test to production.
- Increase business agility by moving to a dynamic platform.
- Streamline test and development efforts with multiple iterations of the same basic image, differencing hard drives, and undo disks.
- Rapidly duplicate a production environment.
- Access operating system and software CD image from virtual media libraries.
- Schedule test environment provisioning.
- Utilize templates to deploy multiple virtual machines at any given time in a single operation.

Standardization

Hardware and operating system standardization has been considered a best practice for many years. IT administrators running datacenters with standardized server hardware and operating system images typically have fewer headaches than those working in mixed environments. This best practice evolved in the physical world and is more significant now when it comes to virtualization and Cloud-based computing.

Standardizing with two or three types of server hardware and one or two operating system images has the advantages of reusable common components and a flexible and adaptable environment, providing a higher level of technical awareness, simplifying upgrade, and easing ongoing management.

Standardization within a virtualized environment is achieved slightly differently. Although the same standards as mentioned earlier apply to the Hyper-V host, now we are dealing with virtual hardware, and this is where virtual machine templates come in.

With templates, you can avoid many repetitive installation and configuration tasks. The result is a fully installed, ready-to-operate virtual machine in less time than a manual installation could ever achieve. Templates are also used to help enforce consistency and standards. Deploying from templates helps enforce corporate standards for such things as hotfixes/patches, hardening, antivirus, and management software in any machine connected to the corporate network.

A virtual machine template is a library resource consisting of the following parts:

- Virtual hard disk
- Hardware profile
- Guest operating system profile

Rapid Deployment

Rapid deployment allows administrators to take advantage of SAN provider technologies to clone a LUN containing a virtual hard disk and present it to the host while still utilizing the SCVMM template so the operating system customization and Integration Services installation can still be applied. This occurs in near real time and removes the need for the virtual hard disk component of a new virtual machine to be copied slowly over the network. This allows and supports a number of differing scenarios:

- Automated and rapid deployment of large virtualized environments
- Automated and rapid deployment of grouped virtual machines concurrently
- Reduced workload for the deployment of similar virtual machines, such as in a VDI scenario
- Provisioning of an environment for disaster recovery and business continuity planning purposes

Greater Flexibility

Virtualization increases an organization's level of flexibility. It removes the dependency between the operating system and the hardware and allows you to grow, shrink, or move your virtual machines, without having to modify the underlying hardware used.

Virtualization allows you to manage your production environment more flexibly, from anywhere, at lower costs and with a reduced level of risk. By leveraging virtualization, you can provide small-scale environments that are cost effective and that scale up very easily.

Hyper-V R2 increases this level of flexibility by introducing a new processor compatibility feature. Processor compatibility allows you to move a running virtual machine to a physical computer with different set of processor features, without having to restart the virtual machine.

This setting may reduce the overall performance of the application in the virtual machine on nodes that would otherwise support advanced virtualization hardware techniques. However, it will allow virtual machines to be live or quick migrated between nodes of differing processor capabilities.

High Availability

Leveraging virtualization for high availability purposes provides businesses with a vast array of high availability solutions from both Microsoft and third parties. These solutions provide high availability for applications that need to recover from failures to complete fault-tolerant solutions for those critical applications that must run continuously without service interruption.

Virtual machines that do not have cluster-aware software can leverage the high availability features in Windows 2008 R2 to implement high availability through host-based clustering. The level of availability that a host-based clustering provides is not as high as with application-specific clustering or guest-based clustering, because the operating systems or applications deployed within the virtual machines are not necessarily cluster-aware.

With Hyper-V R2, a configuration with host-based clustering provides support for both planned and unplanned downtime. During unplanned downtime, the virtual machines will be restarted on a node within the cluster, and during planned downtime, the virtual machines will be transferred, either via Live or Quick Migration, from one node to another.

Conversely, guest-based clustering enables high availability of services and applications at the virtual layer.

The Private Cloud

The private Cloud is an internal service-oriented infrastructure, optimized for both performance and cost, which is deployed inside your datacenter. You can think of the private Cloud as IT-as-a-service, where virtual machines are provisioned to meet business demand. Virtualization effectively unlocks Cloud computing and is a fundamental building block of Cloud computing. The private Cloud in Microsoft's terms is powered by a number of different server products, including Windows 2008 R2, Hyper-V R2, and the System Center family of products. The private Cloud offers a number of benefits:

- A flexible and familiar infrastructure with a common platform to build and deploy applications between clouds and reduce development and deployment time on new services
- Integrated resource access that enables federated services between clouds, helping ensure capacity and resources needed to achieve the business requirements
- The agility to develop applications and services once and then deploy them in and across any cloud environment, enabling rapid response to changing business needs